Likelihood smoothing using gravitational wave surrogate models

Robert H. Cole* and Jonathan R. Gair
Institute of Astronomy, Madingley Road, Cambridge, CB3 0HA, United Kingdom
(Dated: December 16, 2014)

Likelihood surfaces in the parameter space of gravitational wave signals can contain many secondary maxima, which can prevent search algorithms from finding the global peak and correctly mapping the distribution. Traditional schemes to mitigate this problem maintain the number of secondary maxima and thus retain the possibility that the global maximum will remain undiscovered. By contrast, the recently proposed technique of likelihood transform can modify the structure of the likelihood surface to reduce its complexity. We present a practical method to carry out a likelihood transform using a Gaussian smoothing kernel, utilising gravitational wave surrogate models to perform the smoothing operation analytically. We demonstrate the approach with Newtonian and post-Newtonian waveform models for an inspiralling circular compact binary.

arXiv:1410.0597v2 [gr-qc] 23 Jan 2015

I. INTRODUCTION

The first direct detection of gravitational waves (GWs) is likely to occur soon. The ground-based interferometers LIGO [1] and Virgo [2] are currently undergoing upgrades to Advanced configurations which should start taking data in the next couple of years and, when they achieve their final design sensitivity, are expected to detect GWs from the inspiral and merger of stellar compact binaries at the rate of several events per year [3]. There are ongoing efforts to detect a stochastic background of nanohertz GWs generated by merging supermassive black hole binaries using the accurate timing of arrays of millisecond pulsars and the first results could come within five years [4]. Further in the future, the European Space Agency has selected "The Gravitational Universe" to be the science theme for the L3 science mission to launch in 2034, which aims to detect millihertz GWs using a space-based interferometer [5]. These distinct efforts to measure GWs span the wide frequency range of potential sources [6] and it is expected that GW observations from many different sources will eventually become routine [7].

For many astrophysical systems, we can produce accurate models of the source and hence predict the waveform that would be observed on Earth. Given a bank of predicted waveform templates, a GW detection can be made using matched filtering: the comparison of observed data with every template in the bank. For this to be an effective technique, the templates need to be closely separated in parameter space; in high dimensional spaces, it is not possible to construct a template bank of sufficient density using reasonable computational resources. Rather, it is common to use a Markov chain Monte Carlo (MCMC) method to map the *posterior* distribution, evaluating waveforms as required.

The posterior is the probability of a particular set of model parameters given some observed data. To achieve a mapping of the posterior surface with high enough resolution requires many evaluations of the *likelihood* function, which is the probability of producing a particular data stream given some waveform model parameters. The Metropolis-Hastings algorithm provides a technique to explore the parameter space using a random walk, producing a chain of samples that will eventually converge to the posterior distribution. There is no guarantee about how long this convergence will take and the number of likelihood evaluations may still be large. Techniques to accelerate the convergence of MCMC routines are highly desirable.

Simulated annealing [8] is one such method. The idea is to "heat up" the likelihood surface by replacing $\mathcal{L}$ with $\mathcal{L}_T = \mathcal{L}^{1/T}$ for some temperature $T > 1$, making it easier for chains to explore large regions of parameter space. Following some predetermined cooling schedule, the temperature is gradually reduced to $T = 1$, where the chain begins to sample the true likelihood, starting from a point that is more likely to be near the global maximum. This approach is advantageous because it does not slow down the likelihood evaluation at each point, but the number of local maxima of $\mathcal{L}$ remains fixed. Within some fixed computational time, it is possible for a chain to remain close to a secondary and not find the global maximum.

Parallel tempering is related to simulated annealing in that it uses the modified likelihoods $\mathcal{L}_T$. Chains are run simultaneously on a ladder of different temperatures, with high T chains exploring more of the parameter space. Swaps between the locations of adjacent chains are proposed, and in this way information about the global structure of the surface is propagated down to the $T = 1$ chain, which is sampling the desired distribution.

Likelihood transform techniques [9] were recently suggested as an alternative and aim to accelerate MCMC convergence by modifying the likelihood surface in a more complicated way; specifically, we consider the case where $\mathcal{L}$ is convolved with a smoothing kernel $\mathcal{K}_\sigma$. This reduces the number of local maxima, but at the cost of an increased evaluation time at each point.

A separate approach to the problem is to speed up the individual likelihood evaluations by using reduced order methods [10]. These accelerate the likelihood calculation by reducing the number of time or frequency samples at

* rhc26@ast.cam.ac.uk

which the waveforms need to be evaluated by first finding a reduced spanning set for the waveform space. Similarly, the construction of surrogate models [11, 12] achieves acceleration of waveform computations by interpolating the waveform space. In both approaches, significant computational work is done offline to produce the interpolations, allowing for a quicker online run-time.

In this paper, we combine the principles of waveform interpolation and likelihood transformation via a smoothing convolution. A practical scheme for applying likelihood transform methods to GW data analysis with surrogate models is developed, allowing for accelerated convergence in MCMC searches without additional computational time. In section II, we introduce our notation for gravitational waveforms and outline the generation of surrogate models. The principle behind likelihood transform techniques is discussed in section III, along with the application of surrogate models to this problem. We then present some specific examples in section IV, before concluding in section V with a discussion.

II. GRAVITATIONAL WAVES

A gravitational waveform $h(t;\boldsymbol{\lambda})$, depending on some set of parameters $\boldsymbol{\lambda}$, has two independent components, describing two polarisations: plus + and cross ×. A given GW detector is sensitive to a particular linear combination of the waveform polarisations

$$h_\alpha(t;\boldsymbol{\lambda}) = F_\alpha^+ h_+(t;\boldsymbol{\lambda}) + F_\alpha^\times h_\times(t;\boldsymbol{\lambda}), \tag{1}$$

where a subscript α denotes a specific detector and the response functions F_α^A depend on the relative orientations of the detector and the GW source. For the initial generation of ground-based detectors, these are essentially constant over the duration of a typical signal, but they may vary significantly over an observation for space-based interferometers as well as advanced ground-based detectors.

In data analysis, it is necessary to account for the noise present in GW detectors, which is assumed to be stationary and Gaussian. The natural overlap between waveforms is in the frequency domain

$$\langle\, h(\boldsymbol{\lambda}_1) \,|\, h(\boldsymbol{\lambda}_2) \,\rangle \equiv 4 \sum_\alpha \int_0^\infty \frac{\tilde{h}_\alpha^*(f;\boldsymbol{\lambda}_1)\tilde{h}_\alpha(f;\boldsymbol{\lambda}_2)}{S_{n,\alpha}(f)}\,\mathrm{d}f, \tag{2}$$

where we sum over different detectors, each with their own one-sided noise power spectral density $S_{n,\alpha}(f)$. Here $\tilde{h}_\alpha(f)$ denotes the Fourier transform of $h_\alpha(t)$, and we are using $\langle\,\cdot\,|\,\cdot\,\rangle$ to denote a noise-weighted overlap.

It is often convenient to write the waveform as a complex time series

$$h(t;\boldsymbol{\lambda}) = h_+(t;\boldsymbol{\lambda}) + i h_\times(t;\boldsymbol{\lambda}). \tag{3}$$

The natural inner product on this complex waveform space is

$$(\, h(\boldsymbol{\lambda}_1) \,|\, h(\boldsymbol{\lambda}_2) \,) \equiv \int_{-\infty}^{\infty} h^*(t;\boldsymbol{\lambda}_1) h(t;\boldsymbol{\lambda}_2)\,\mathrm{d}t, \tag{4}$$

where in practice, the integral is of finite length, T, equal to the observation time. The corresponding real overlap between two waveforms is given by the real part of this inner product. The complex inner product and associated overlap make no reference to a particular detector and are therefore useful for constructing reduced spanning sets for waveform spaces [12]. The two overlaps coincide if it is assumed that there are two right-angled detectors, at 45° to one another, with independent white noise, $S_{n,1} = S_{n,2} = \text{const.}$, and that the source is optimally oriented, i.e., the principal polarisation axes of the source are aligned with the arms of the first detector. In this configuration the detector aligned with the principal axes is sensitive only to the plus polarisation, while the other detector is sensitive only to the cross polarisation and the overlap is proportional to

$$\sum_{A=+,\times} \int_{-\infty}^{\infty} h_A(t;\boldsymbol{\lambda}_1) h_A(t;\boldsymbol{\lambda}_2)\,\mathrm{d}t = \Re[\,(\, h(\boldsymbol{\lambda}_1) \,|\, h(\boldsymbol{\lambda}_2) \,)\,], \tag{5}$$

where the sum is over polarisation states. We will assume this optimal configuration for all sources in subsequent calculations, as the likelihood transform technique is independent of our choice of detector.

A. Data analysis

Given some measured detector data $x_\alpha(t) = h_\alpha(t;\boldsymbol{\lambda}_*) + n_\alpha(t)$, composed of a GW signal h with parameters $\boldsymbol{\lambda}_*$ and detector noise n, the signal-to-noise ratio (SNR) can be calculated

$$\rho(\boldsymbol{\lambda}) = \frac{\langle\, x \,|\, h(\boldsymbol{\lambda}) \,\rangle}{\sqrt{\langle\, h(\boldsymbol{\lambda}) \,|\, h(\boldsymbol{\lambda}) \,\rangle}}, \tag{6}$$

which we expect to be strongly peaked at the true parameters $\boldsymbol{\lambda}_*$. It is useful to work with normalised templates, such that $\langle\, h(\boldsymbol{\lambda}) \,|\, h(\boldsymbol{\lambda}) \,\rangle = 1$. For noise-free data, $x_\alpha(t) = h_\alpha(t;\boldsymbol{\lambda}_*)$, and making the simplifying assumptions about the source orientation and noise properties described above, the SNR reduces to

$$\rho(\boldsymbol{\lambda};\boldsymbol{\lambda}_*) = \Re[\,(\, h(\boldsymbol{\lambda}_*) \,|\, h(\boldsymbol{\lambda}) \,)\,], \tag{7}$$

which is linear in the model $h(\boldsymbol{\lambda})$.

The likelihood $\mathcal{L}(x|\boldsymbol{\lambda})$ is the probability that a particular data stream x is observed, given the parameters $\boldsymbol{\lambda}$ of the signal present. Assuming stationary Gaussian noise, the likelihood is simply

$$\mathcal{L}(x|\boldsymbol{\lambda}) \propto \exp\left[-\,\langle\, x - h(\boldsymbol{\lambda}) \,|\, x - h(\boldsymbol{\lambda}) \,\rangle\,/\,2\right], \tag{8}$$

up to some normalising factor. The actual quantity of interest is the posterior $\mathcal{P}(\boldsymbol{\lambda}|x)$, which is the probability

distribution of the parameters of the source $\boldsymbol{\lambda}$, given the observed data stream x. It is related to $\mathcal{L}$ via Bayes' theorem

$$\mathcal{P}(\boldsymbol{\lambda}|x) = \frac{\mathcal{L}(x|\boldsymbol{\lambda})\pi(\boldsymbol{\lambda})}{Z}, \qquad Z = \int \mathrm{d}\boldsymbol{\lambda}\, \mathcal{L}(x|\boldsymbol{\lambda})\pi(\boldsymbol{\lambda}) \quad (9)$$

where $\pi(\boldsymbol{\lambda})$ is the prior probability distribution for the parameters, which reflects our beliefs about the source parameters prior to the data being taken. The evidence, Z, normalises the posterior over the parameter space and can also be used for model selection.

The expectation value of the likelihood over different noise realisations can be expanded about the true parameters in the linear signal approximation to yield the Fisher Information Matrix (FIM)

$$\Gamma_{ij} = \langle\, h_{,i} \,|\, h_{,j} \,\rangle\ , \quad (10)$$

where $h_{,i} = \partial h/\partial\lambda_i$ denotes the partial derivative of the waveform model with respect to the parameters. In the case of uninformative (uniform) priors, the inverse of the FIM is the covariance of the posterior distribution, and so can be used to set a scale on its structure. The component $(\Gamma^{-1})_{ii}$ is a measure of the expected width of the marginalised posterior in λ_i and hence is a measure of the expected uncertainty in the measurement of that parameter from the observed data.

The posterior encodes all of the information about the source that can be determined from an observation. In low-dimensional parameter spaces, the posterior can be evaluated on a fine grid in parameter space. For higher dimensional spaces, the required computational time is too large and alternative methods are typically adopted, such as Markov chain Monte Carlo methods.

MCMC techniques aim to generate a chain of samples, $\boldsymbol{\lambda}_i$ in which, after a burn-in phase, the density of points is proportional to the posterior distribution. This is typically achieved using the Metropolis-Hastings algorithm. Given some current parameter value $\boldsymbol{\lambda}_i$, a new candidate parameter value $\boldsymbol{\lambda}'$ is chosen from a suitable proposal distribution $q(\boldsymbol{\lambda}'|\boldsymbol{\lambda}_i)$, for instance a Gaussian centred at $\boldsymbol{\lambda}_i$. The Metropolis-Hastings ratio

$$\alpha = \frac{\mathcal{P}(\boldsymbol{\lambda}'|x)q(\boldsymbol{\lambda}_i|\boldsymbol{\lambda}')}{\mathcal{P}(\boldsymbol{\lambda}_i|x)q(\boldsymbol{\lambda}'|\boldsymbol{\lambda}_i)} \quad (11)$$

is calculated and the new state is accepted with probability $\min(1, \alpha)$. Otherwise, the next state is set to be the current parameter values $\boldsymbol{\lambda}_{i+1} = \boldsymbol{\lambda}_i$. The starting point for the algorithm, $\boldsymbol{\lambda}_0$, can be chosen arbitrarily.

B. Reduced order methods and surrogate models

Gravitational waveforms are routinely generated for arbitrary system parameters by numerically solving differential equations. The accuracy of such waveforms is not guaranteed (and in some cases, deliberately sacrificed to reduce the computational cost, for example kludge models of extreme-mass-ratio inspirals [13, 14]); indeed, it is possible to get significant systematic biases in parameter estimation by using unfaithful waveform templates, but we shall assume in this analysis that the numerical templates can be calculated exactly. Waveform models and the associated likelihood, Eq. (8), can be expensive to evaluate and reduced order methods have been proposed as a way to speed up such likelihood evaluations. These rely on finding an approximation to the likelihood, employing a surrogate model for the waveform, that is cheaper to evaluate.

There are many possible ways of approaching this problem [11, 12]. Here, we give a brief summary of a procedure for generating a reduced order likelihood and surrogate waveform model. This description follows Field *et al* [12], where more details may be found.

As a starting point, we compute M waveform templates $h(t; \boldsymbol{\lambda}_i)$, referred to as the *training set*. The aim is first to find the minimal number m of these waveforms such that all other waveforms in the training set can be well approximated. If the training set is sufficiently dense, the approximation will also be valid for waveforms outside the training set. Once this reduced basis set has been found, the second stage is to identify a set of m discrete times at which it is sufficient to compute the waveform in order to represent it faithfully with the reduced basis. The final stage is to construct a surrogate model that can predict the value of the waveform at those times for arbitrary choices of the model parameters. The algorithm for achieving this is as follows:

1. Choose the m most differing waveforms and construct an orthonormal basis $\{e_i(t)\}_{i=1}^{i=m}$ from them[1]. The corresponding waveform parameters are referred to as *greedy data*. Given the similarity between GWs with different parameters, it is expected that $m \ll M$. Every waveform in the training set, as well as waveforms that are not in the training set, may then be approximated by the expansion

$$h_m(t; \boldsymbol{\lambda}) \approx \sum_{i=1}^{m} c_i(\boldsymbol{\lambda}) e_i(t). \quad (12)$$

The representation error of the reduced basis (RB) is defined as

$$\sigma_m \equiv \max_{\boldsymbol{\lambda}} \min_{c_i} \left(h(\boldsymbol{\lambda}) - h_m(\boldsymbol{\lambda}) \,|\, h(\boldsymbol{\lambda}) - h_m(\boldsymbol{\lambda}) \right)\ , \quad (13)$$

where the minimisation over the coefficients is done by projecting the waveform onto the RB. The value of m is set by a condition on σ_m; typically we may require that $\sigma_m \lesssim \mathcal{O}(10^{-12})$.

[1] The basis is constructed in a greedy manner. Given an existing basis of size $r < m$, the $(r+1)$th waveform is the member of the training set with the largest residual norm when projected onto the r-basis. The first waveform is chosen arbitrarily.

2. Identify m evaluation times $\{T_i\}_{i=1}^{i=m}$ that can be used to construct an empirical interpolant for the RB; these are referred to as *empirical nodes*[2]. The goal is to construct an interpolant

$$\mathcal{I}_m[h](t;\boldsymbol{\lambda}) = \sum_{j=1}^{m} B_j(t) h(T_j;\boldsymbol{\lambda}), \tag{14}$$

where

$$B_j(t) \equiv \sum_{i=1}^{m} e_i(t)(V^{-1})_{ij} \tag{15}$$

is independent of the system parameters and so may be computed offline. The V matrix is constructed by requiring that $\mathcal{I}_m[h](T_j;\boldsymbol{\lambda}) = h(T_j;\boldsymbol{\lambda})$.

3. At each empirical node, predict the waveform value for arbitrary parameters by fitting $h(T_i;\boldsymbol{\lambda})$ with respect to $\boldsymbol{\lambda}$, using only the greedy data. It is often easier to find fits for the amplitude A_i and phase ϕ_i of the waveform independently (as opposed to h_+ and h_x individually). The waveform can then be written as

$$h(T_i;\boldsymbol{\lambda}) \approx A_i(\boldsymbol{\lambda}) e^{i\phi_i(\boldsymbol{\lambda})}, \tag{16}$$

where the fits are arbitrary functions of the parameters; the specific choice will be determined by the problem at hand.

Once these offline steps have been completed, a surrogate model can be constructed from the empirical interpolant

$$h_S(t;\boldsymbol{\lambda}) = \sum_{i=1}^{m} B_i(t) A_i(\boldsymbol{\lambda}) e^{i\phi_i(\boldsymbol{\lambda})}. \tag{17}$$

The model is quick to evaluate at arbitrary $\boldsymbol{\lambda}$ as the $\{B_i\}$ are computed in advance, and the $\{A_i\}$ and $\{\phi_i\}$ are simple analytic expressions, fitted to the numerical data.

III. LIKELIHOOD TRANSFORM

A general optimisation problem involves finding the set of parameters that globally maximises some function. Here we will initially calculate the SNR to illustrate the smoothing technique in a detector- and noise-independent way. We then apply the method to the maximisation of a simple noise-dependent likelihood function.

[2] The empirical nodes are selected in a greedy manner. Using r existing nodes and the first r basis functions, an empirical interpolant (14) is built for the $(r+1)$th basis element. The $(r+1)$th empirical node is the time at which this interpolant is most different from the actual basis function.

The SNR often has a great deal of structure: there are many local secondary maxima and the true global peak is tall and narrow. This makes the optimisation problem difficult because a high resolution is required and it is possible that any algorithm will stall at a secondary maximum, and thus not explore the full parameter space. It is therefore desirable to be able to sample a smoother distribution that closely mirrors the true SNR surface. To this extent, we follow Wang [9] and define a smoothed SNR via a convolution

$$\rho_\sigma(\boldsymbol{\lambda};\boldsymbol{\lambda}_*) \equiv \mathcal{K}_\sigma \star \rho(\boldsymbol{\lambda};\boldsymbol{\lambda}_*) = \Re[\,(\,h(\boldsymbol{\lambda}_*)\,|\,h_\sigma(\boldsymbol{\lambda})\,)\,], \tag{18}$$

where $\mathcal{K}_\sigma$ is some smoothing kernel of typical width σ, we have used the linear property of the SNR and

$$h_\sigma(\boldsymbol{\lambda}) \equiv \mathcal{K}_\sigma \star h(\boldsymbol{\lambda}) = \int \mathcal{K}_\sigma(\boldsymbol{\lambda}-\boldsymbol{\lambda}') h(t;\boldsymbol{\lambda}')\,\mathrm{d}\boldsymbol{\lambda}'. \tag{19}$$

It would be possible to numerically perform this integral, by sampling the waveform space at k points surrounding the desired parameters $\boldsymbol{\lambda}$ and compute the appropriate sum, weighted by $\mathcal{K}_\sigma$. The major drawback to this approach is the required computational time; running time would be increased by $\mathcal{O}(k)$, and in high N-dimensional parameter spaces, the required number of points to accurately estimate the integral may be very large, $k \sim \mathcal{O}(2^N)$.

Alternatively, if there exists an analytic expression for the waveform model $h(t;\boldsymbol{\lambda})$, it may be possible to explicitly calculate (or at least approximate) the integral for particular choices of kernel $\mathcal{K}_\sigma$. This would result in an analytic expression for h_σ, enabling the smoothed SNR to be calculated quickly. This was the approach taken in the previous study of likelihood transform methods [9], where a simple quadratic chirp signal was considered. This approach is highly restrictive since faithful models of likely GW sources are not usually analytic but instead are computed numerically on some parameter grid. The quadratic chirp model used in [9] is not a faithful model of any likely GW signal.

We take an approach between these two extremes, showing how the surrogate model (17) can be utilised to simplify the smoothing operation (19) required in likelihood transform techniques[3]. To do this, we must first lose some generality, although the resulting procedure remains general enough to be of wide applicability.

The N as-yet unspecified parameters, $\boldsymbol{\lambda}$, for each waveform are mapped onto the unit cube $\boldsymbol{\theta}$ in parameter space. This allows us to discuss a wide range of waveform models, without being overly specific. We could have equally chosen to map the parameters $\boldsymbol{\lambda}$ onto an infinite or semi-infinite range, giving similar results.

[3] We later illustrate the technique using an analytic model but, in contrast to Wang [9], our procedure does not rely on this.

We choose the kernel $\mathcal{K}_\sigma$ to be a multivariate Gaussian with diagonal covariance matrix $\mathbf{\Sigma} = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_N^2)$. The smoothed waveform can then be calculated using

$$h_\sigma(\boldsymbol{\theta}) = \int \mathrm{d}\boldsymbol{\theta}'\, h(t;\boldsymbol{\theta}') \prod_{j=1}^{N} \mathcal{N}_j e^{-(\theta_j-\theta_j')^2/2\sigma_j^2}, \tag{20}$$

where $\mathcal{N}_j \equiv \mathcal{N}(\theta_j;\sigma_j)$ is a normalisation function that we derive later.

We now make use of the surrogate model (17) to remove the time dependence from the integrals

$$h_\sigma(\boldsymbol{\theta}) = \sum_{i=1}^{m} B_i(t) \int \mathrm{d}\boldsymbol{\theta}'\, A_i(\boldsymbol{\theta}') e^{i\phi_i(\boldsymbol{\theta}')} \prod_{j=1}^{N} \mathcal{N}_j e^{-(\theta_j-\theta_j')^2/2\sigma_j^2}. \tag{21}$$

The phase of the waveform may be approximated around the evaluation point using a Taylor series

$$\phi_i(\boldsymbol{\theta}') \approx \phi_i(\boldsymbol{\theta}) + (\boldsymbol{\theta}' - \boldsymbol{\theta}).\nabla\phi_i. \tag{22}$$

This is a good approximation when $(\boldsymbol{\theta}' - \boldsymbol{\theta})$ is small, which is true if we choose σ to be sufficiently small[4]. The waveform can then be written as

$$h_\sigma(\boldsymbol{\theta}) = \sum_{i=1}^{m} B_i(t) e^{i\phi_i(\boldsymbol{\theta})} \int \mathrm{d}\boldsymbol{\theta}'\, A_i(\boldsymbol{\theta}') \prod_{j=1}^{N} \mathcal{N}_j e^{-(\theta_j'-\theta_j)^2/2\sigma_j^2} \exp\left[i(\theta_j'-\theta_j)\frac{\partial\phi_i}{\partial\theta_j}\right]. \tag{23}$$

We are free to choose any set of functions $\{A_i\}$, as long as they accurately fit the numerical waveform data. We first focus on functions that can be decomposed into a short series of separable terms[5]

$$A_i(\boldsymbol{\theta}) = \sum_{\text{terms}} \prod_{j=1}^{N} f_j(\theta_j), \tag{24}$$

where f_j are arbitrary functions that may in principle be different for each term. We will discuss generic amplitudes later. The smoothed waveform can then be written as

$$h_\sigma(\boldsymbol{\theta}) = \sum_{i=1}^{m} B_i(t) e^{i\phi_i(\boldsymbol{\theta})} \sum_{\text{terms}} \int \mathrm{d}\boldsymbol{\theta}' \prod_{j=1}^{N} \mathcal{N}_j f_j(\theta_j')\, e^{-(\theta_j'-\theta_j)^2/2\sigma_j^2} \exp\left[i(\theta_j'-\theta_j)\frac{\partial\phi_i}{\partial\theta_j}\right], \tag{25}$$

which can be simplified to

$$h_\sigma(\boldsymbol{\theta}) = \sum_{i=1}^{m} B_i(t) e^{i\phi_i(\boldsymbol{\theta})} \sum_{\text{terms}} \prod_{j=1}^{N} f_j(\theta_j;\sigma_j), \tag{26}$$

where we define

$$f_j(\theta_j;\sigma_j) \equiv \int_0^1 \mathrm{d}\theta_j'\, \mathcal{N}_j\, f_j(\theta_j')\, e^{-(\theta_j'-\theta_j)^2/2\sigma_j^2} \exp\left[i(\theta_j'-\theta_j)\frac{\partial\phi_i}{\partial\theta_j}\right]. \tag{27}$$

The appeal of such an approach is immediate: the smoothed waveform takes an identical form to the surrogate model, but with the replacement $f_j(\theta_j) \to f_j(\theta_j;\sigma_j)$.

A. Polynomial amplitudes

We now consider the specific case of amplitude functions $A_i(\boldsymbol{\theta})$ that can be well described by polynomials. In this case, all of the $f_j(\theta_j)$ will be powers of θ_j. We are hence interested in integrals of the form

$$f(\theta;\sigma,n) \equiv \int_0^1 \mathrm{d}\theta' \mathcal{N} \theta'^n\, e^{-(\theta'-\theta)^2/2\sigma^2} \exp\left[i(\theta'-\theta)\frac{\partial\phi_i}{\partial\theta}\right], \tag{28}$$

where we have dropped the j subscripts, for clarity. Completing the square in the exponential terms leads to

$$f(\theta;\sigma,n) = \mathcal{N} \exp\left[-\frac{\sigma^2}{2}\left(\frac{\partial\phi_i}{\partial\theta}\right)^2\right] C^n(\theta + i\sigma^2\frac{\partial\phi_i}{\partial\theta};\sigma), \tag{29}$$

where we have defined

$$C^n(z;\sigma) \equiv \int_0^1 \mathrm{d}z'\, z'^n\, e^{-(z'-z)^2/2\sigma^2}. \tag{30}$$

Integrating this by parts gives the recurrence relation

[4] Even for larger values of σ, the procedure can still be followed. In this case, the smoothed waveform (23) is not a good approx-

$$C^n(z;\sigma) = \frac{1}{(n+1)} e^{-(1-z)^2/2\sigma^2} + \frac{1}{\sigma^2(n+1)} \int_0^1 \mathrm{d}z' \left\{ z'^{n+2} - z z'^{n+1} \right\} e^{-(z'-z)^2/2\sigma^2}, \tag{31}$$

$$= \frac{1}{(n+1)} e^{-(1-z)^2/2\sigma^2} + \frac{1}{\sigma^2(n+1)} \left\{ C^{n+2}(z;\sigma) - z C^{n+1}(z;\sigma) \right\}, \tag{32}$$

$$= z C^{n-1}(z;\sigma) + \sigma^2 \left((n-1) C^{n-2}(z;\sigma) - e^{-(1-z)^2/2\sigma^2} \right), \tag{33}$$

where the last line follows from relabelling n by $n-2$ in (32). We can perform the integral explicitly for the first two terms:

$$C^0(z;\sigma) = \sqrt{\frac{\pi}{2}} \sigma \left(\mathrm{erf}\left(\frac{1-z}{\sqrt{2}\sigma} \right) + \mathrm{erf}\left(\frac{z}{\sqrt{2}\sigma} \right) \right), \tag{34}$$

$$C^1(z;\sigma) = z C^0(z;\sigma) + \sigma^2 \left(e^{-z^2/2\sigma^2} - e^{-(1-z)^2/2\sigma^2} \right). \tag{35}$$

To correctly normalise the Gaussian kernel, we must set $\mathcal{N}_j = 1/C^0(\theta_j;\sigma_j)$. To construct a smoothed waveform, we write out the surrogate model, including the polynomial fit for the amplitude functions, and then make the replacements[6]

$$\theta_j^n \to f(\theta_j;\sigma_j,n). \tag{36}$$

We note the expected property that $f(\theta;\sigma=0,n) = \theta^n$, thus recovering the original polynomial in the case of zero smoothing.

B. Fourier amplitudes

Rather than polynomials, it may be desirable to decompose the amplitude functions into Fourier components

$$A(\boldsymbol{\theta}) = \sum_{k_1,k_2,\ldots k_N} A_{\boldsymbol{k}} e^{2\pi i \boldsymbol{k}.\boldsymbol{\theta}}, \tag{37}$$

where we have dropped the i subscript for clarity, $\boldsymbol{k} = \{k_1, k_2, \ldots k_N\}$, the sum runs over positive and negative integers, and

$$A_{\boldsymbol{k}} = \int_0^1 A(\boldsymbol{\theta}) e^{-2\pi i \boldsymbol{k}.\boldsymbol{\theta}} \, \mathrm{d}\boldsymbol{\theta}. \tag{38}$$

We note that (37) is of the form of (24) and so we are interested in computing integrals of the form

$$f(\theta;\sigma,n,k) \equiv \int_0^1 \mathrm{d}\theta' \, \mathcal{N} \, \theta'^n \, e^{ik\theta'} \, e^{-(\theta'-\theta)^2/2\sigma^2} \exp\left[i(\theta'-\theta) \frac{\partial \phi_i}{\partial \theta} \right], \tag{39}$$

where we have also included a polynomial factor, for generality. Following the same procedure as before leads to

$$f(\theta;\sigma,n,k) = \mathcal{N} \exp\left[-\frac{\sigma^2}{2} \left(\frac{\partial \phi_i}{\partial \theta} + k \right)^2 + ik\theta \right] C^n\left(\theta + i\sigma^2 \left(\frac{\partial \phi_i}{\partial \theta} + k \right); \sigma \right), \tag{40}$$

where the calculation of C^n is discussed in the previous section. These can then be used in (37), replacing each exponential factor according to

$$e^{2\pi i k \theta} \to f(\theta;\sigma,0,2\pi k). \tag{41}$$

C. Generic amplitudes

In some cases, it may not be possible to decompose the amplitude functions into a series of appropriate separable functions. In this case, we may use a Taylor series to approximate

$$A_i(\boldsymbol{\theta}') \approx A_i(\boldsymbol{\theta}) + (\boldsymbol{\theta}' - \boldsymbol{\theta}).\nabla A_i. \tag{42}$$

As with the Taylor series in the phase (22), this is a good approximation for sufficiently small smoothing widths σ. Substituting it into (23) results in

$$h_\sigma(\boldsymbol{\theta}) = \sum_{i=1}^{m} B_i(t) e^{i\phi_i(\boldsymbol{\theta})} \Bigg\{ A_i(\boldsymbol{\theta}) \prod_{j=1}^{N} f(\theta_j;\sigma_j,0) + \sum_{k=1}^{N} \int_0^1 \mathrm{d}\theta_k' \mathcal{N}_k \left(\theta_k' - \theta_k\right) \frac{\partial A_i}{\partial \theta_k} e^{-(\theta_k'-\theta_k)^2/2\sigma_k^2} \exp\left[i(\theta_k' - \theta_k) \frac{\partial \phi_i}{\partial \theta_k} \right] \prod_{j\neq k} f(\theta_j;\sigma_j,0) \Bigg\}, \quad (43)$$

which can be written as

$$h_\sigma(\boldsymbol{\theta}) = \sum_{i=1}^{m} B_i(t) e^{i\phi_i(\boldsymbol{\theta})} \left\{ A_i(\boldsymbol{\theta}) + \sum_{k=1}^{N} \frac{\partial A_i}{\partial \theta_k} \frac{f(\theta_k;\sigma_k,1) - \theta_k f(\theta_k;\sigma_k,0)}{f(\theta_k;\sigma_k,0)} \right\} \prod_{j=1}^{N} f(\theta_j;\sigma_j,0), \quad (44)$$

where $f(\theta;\sigma,n)$ is defined in (29).

With these results, it should be possible to produce a smoothed SNR or likelihood for any given surrogate model waveform. We note that the smoothed waveforms do not necessarily have to be accurate (neglected terms of higher order in A_i and ϕ_i may be large) as long as the resulting smoothed surface displays the desired properties that the global maximum is broadened and the number of secondary maxima has been reduced.

Using (44), the expected additional computational cost of evaluating a smoothed waveform with N parameters is $\mathcal{O}(N)$, in comparison to $\mathcal{O}(2^N)$ for naive likelihood smoothing.

IV. CHIRP WAVEFORMS

As an illustration of our method, we consider the gravitational waveforms expected from a circular compact binary. Such waveforms are well-approximated by a high-order post-Newtonian (PN) expansion [15] in the dimensionless variable $x = (GM\Omega/c^3)^{2/3}$, where M is the total mass of the system and $\Omega = \dot{\Phi}$ is the angular frequency of the binary, computed as a time derivative of the binary phase Φ. The expansion takes the form

$$h_{+,\times} = \frac{2G\mu x}{c^2 R} \sum_{p=0}^{\infty} x^{p/2} \underset{p/2}{H}_{+,\times}(\psi,\iota;\ln x) + \mathcal{O}(R^{-2}), \quad (45)$$

where μ is the reduced mass, R is the distance to the source, ι is the inclination of the binary and ψ is the binary phase distorted by tails

$$\psi = \Phi - \frac{2GM_{\mathrm{ADM}}}{c^3} \Omega \ln\left(\frac{\Omega}{\Omega_0}\right). \quad (46)$$

The tail integrals are a result of the nonlinear interaction between the source and the emitted GWs [16]. Here one must use the binary's mass monopole M_{ADM}, which includes all contributions to the mass-energy of the binary. At 1PN order, it may be computed as [17]

$$M_{\mathrm{ADM}} = M\left(1 - \frac{\nu}{2}x\right), \quad (47)$$

where $\nu \equiv \mu/M$ is the symmetric mass ratio, which takes values between 0 (test particle limit) and 1/4 (equal mass). Ω_0 is often chosen to be the lower cutoff of the detector band; we make the choice $\Omega_0 = 10\pi\,\mathrm{rad\,s^{-1}}$.

To demonstrate the approach, we now make some simplifying assumptions. First, the expansion will be truncated at finite PN order: we specify the phase Φ up to 3.5PN (this is given explicitly in appendix A); the amplitude to 2PN; and M_{ADM} to 1PN[7]. The logarithmic term in ψ is then at 4PN order relative to the dominant phase contribution, and is included for completeness. Secondly, we will consider aligned binaries that have $\iota = 0$. In this case, the relevant expansion functions for the plus polarisation are all proportional to $\cos(2\psi)$ and the cross terms are proportional to $\sin(2\psi)$. We can therefore write a two-parameter family of complex PN waveforms as

$$h(\mathcal{M},\nu;t) = \frac{2G\mathcal{M}\nu^{2/5}}{c^2 R} x(t) H(x(t),\nu) e^{i(2\psi(t)+\pi)}, \quad (48)$$

where $\mathcal{M} = \mu^{3/5}M^{2/5}$ is the chirp mass and the amplitude function H takes the form [18]

$$H(x,\nu) = 2 + \frac{1}{3}(\nu - 13)\,x + 4\pi x^{3/2} + \frac{1}{180}\left(15\nu^2 - 635\nu - 837\right)x^2. \quad (49)$$

imation to (21), but the resulting smoothed likelihood surface may still be sufficiently similar to the unsmoothed surface that MCMC convergence will be accelerated.

[5] If we allow the number of terms to approach infinity, this decomposition can be used to represent any sufficiently smooth function, but the number of required fitting parameters will also approach infinity. By *short* series, we mean that the number of fitting parameters required to represent the function is smaller than the number of elements m.

[6] In particular, any term that does not depend on θ_j should be multiplied by $f(\theta_j;\sigma_j,0)$.

[7] It is not a requirement to have consistency in PN orders between the amplitude and phase.

The clean separation into an amplitude and phase is a consequence of our simplifying assumptions, but this is not a necessary requirement for the smoothing procedure.

A. Newtonian chirp

We first simplify the waveform model even more, truncating at 0PN order so that the waveform has only one free intrinsic parameter, the chirp mass $\mathcal{M}$. We choose $t_c = 0.4$s, and set a fiducial distance of $R = 1$Mpc, although this is unimportant as we normalise the waveforms to unity. We shall consider chirp masses in the range $1 \leq \mathcal{M}/M_\odot \leq 20$, corresponding roughly to the range considered in recent LIGO searches [19], and inject a signal with $\mathcal{M}_* = 2.2M_\odot$.

We can calculate the FIM for this waveform model at our injection parameters. Using the unit parameter $\theta = (\mathcal{M}/M_\odot - 1)/19$, we find that the global maximum of the likelihood has a characteristic scale of 7×10^{-4}.

1. *Idealised data*

We initially perform calculations without noise, to demonstrate the underlying properties of the smoothing technique. Our injected waveform is plotted in figure 1.

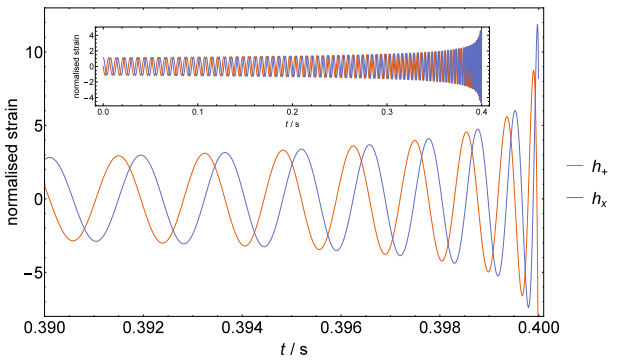

FIG. 1. Gravitational waveform from a compact binary with chirp mass $\mathcal{M} = 2.2M_\odot$ in the final 0.01s before coalescence. Inset is the same waveform for a period of 0.4s before coalescence. The amplitude is set by the requirement that $(h\,|\,h) = 1$.

We construct a surrogate model for our waveform family, as discussed in section II B, using a training set of 501 waveforms, sampled at a frequency $f_{\rm sample} = 20\,$kHz and with chirp masses selected such that they sample the frequency Ω uniformly[8]. We target a representation error of 10^{-12} and the resulting RB contains 133 elements. The error as a function of the size of the RB is plotted in figure 2.

To check the faithfulness of the RB, we compute the representation error (13) for 1000 waveforms *not* in the training set, generated with random values of $\mathcal{M}$ sampled uniformly from the allowed range; all of the waveforms had $\sigma < 10^{-12}$, with typical values of $\sigma \sim 10^{-15}$.

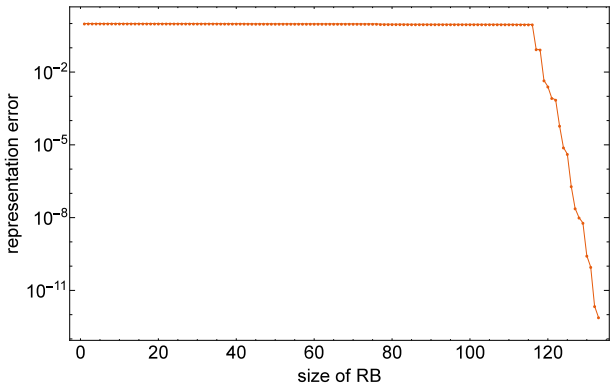

FIG. 2. The representation error of our RB for the 0PN waveform family as a function of the number of elements in the basis.

Using the training set waveforms identified in generating the RB, we perform a fit to the waveform phase at each empirical node, using the functional form

$$\phi_i(\mathcal{M}) = a_1 + a_2\mathcal{M}^{a_3}, \tag{50}$$

where $\{a_k\}_{k=1}^3$ are the fitting parameters. Figure 3 shows the phase as a function of chirp mass, evaluated at the first empirical node, along with the fit; the good agreement is not surprising as the fitting function was motivated by our knowledge of the exact waveform model. Similar results are found for the phase at the other empirical nodes.

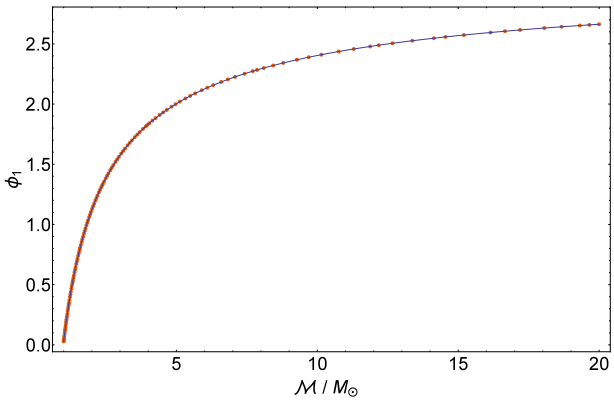

FIG. 3. The phase of the gravitational waveform, evaluated at the first empirical node, as a function of chirp mass. The dots show the values of the phase evaluated at the greedy points selected while constructing the RB. The line is a fit to the data, which in this case is exact.

[8] Selecting a training set with uniform values of $\mathcal{M}$ resulted in a poor surrogate model for low values of $\mathcal{M}$ (high frequencies).

For the 0PN waveforms, the amplitude at each empirical node is simply a constant

$$A_i(\mathcal{M}) = b_i. \tag{51}$$

Following the procedure in section III A, we then make the replacement at each empirical node

$$b_i \to b_i \mathcal{N} \exp\left[-\frac{\sigma^2}{2}\left(\frac{\partial\phi_i}{\partial\theta}\right)^2\right] C^0(\theta + i\sigma^2\frac{\partial\phi_i}{\partial\theta};\sigma), \tag{52}$$

where $\theta = (\mathcal{M}/M_\odot - 1)/19$ is the chirp mass mapped onto the unit cube. The resulting smoothed waveforms for our injection signal for different values of σ are shown in figure 4.

The choice of σ is arbitrary, but clearly has a large impact on the resulting performance of the algorithm: choose σ too small and the smoothed waveforms will be indistinguishable from the unsmoothed waveforms; choose σ too large and the smoothed waveforms will be sufficiently dissimilar that the transform technique will not aid in locating the global maximum. To get some idea of the required scale, we compute the smoothed SNR $\rho_\sigma(\mathcal{M}_*;\mathcal{M}_*)$ of our injection waveform, as a function of smoothing width, shown in figure 5. We see that choosing values $\sigma \sim \mathcal{O}(10^{-3})$ results in smoothed waveforms that are different to the unsmoothed waveform, but which still give a high SNR $\rho_\sigma \gtrsim 0.1$. This is consistent with the typical scale of the peak of the likelihood distribution obtained from the FIM.

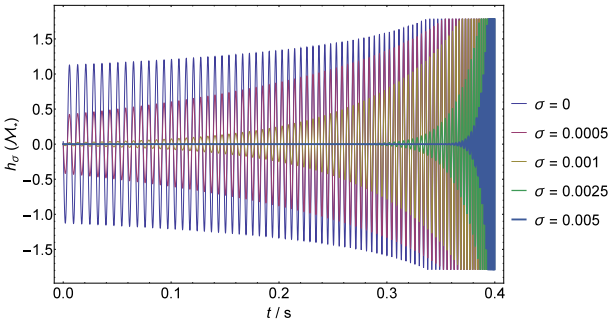

FIG. 4. The smoothed waveform (26), computed for our injection parameters with a chirp mass $\mathcal{M} = 2.2M_\odot$. Thicker lines correspond to larger values of the smoothing width σ. The $\sigma = 0$ waveform is identical to that in figure 1.

To simulate searching for the global maximum, we compute the smoothed SNR $\rho_\sigma(\mathcal{M};\mathcal{M}_*)$ from (18) on a grid of $\{\sigma, \mathcal{M}\}$ values; the resulting curves for different values of σ are shown in figure 6. It can be seen that the desired smoothing properties have been achieved: the global peak has been broadened and the number of secondary maxima has been reduced.

If the likelihood transform is to be useful, the time taken to evaluate ρ_σ must not greatly exceed that taken to evaluate ρ_0[9]. In this example, we found that to compute the smoothed SNR took roughly a factor of 4 longer than ρ_0. The time taken to calculate the training set data and to produce the surrogate model is relatively large, but this can be done offline.

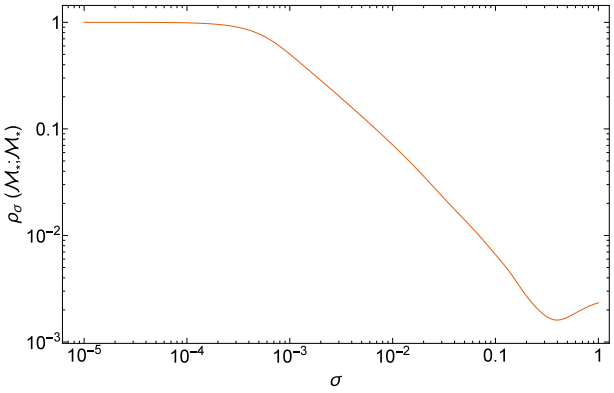

FIG. 5. The smoothed SNR (18), computed for our injection parameters with a chirp mass $\mathcal{M}_* = 2.2M_\odot$, as a function of smoothing width.

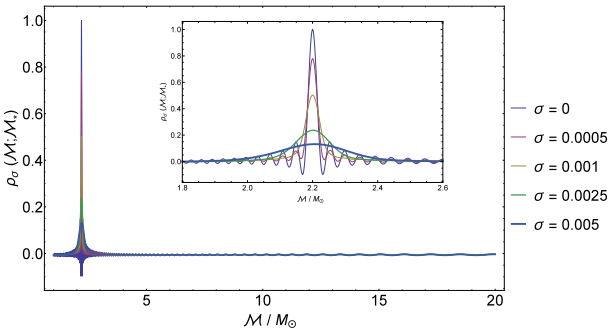

FIG. 6. The smoothed SNRs (18), computed across the allowed range of chirp masses, for a selection of smoothing widths. The inset plot shows the behaviour around the injection value $\mathcal{M}_* = 2.2M_\odot$. Thicker lines correspond to larger values of σ.

2. *Noisy data*

We now consider analysing a data set containing an injected signal and Gaussian white noise. We use the plus polarisation of the waveform only, which is equivalent to having an optimally oriented source observed with a right-angle detector aligned with the principal polarisation axes of the system, as discussed in Section II. The SNR of the injected signal is approximately 400. This is a particularly large SNR, albeit not unusual for, say, supermassive black hole mergers observed with space-based

[9] We compare to ρ_0 rather than ρ to make use of the pre-calculated surrogate model.

detectors. We deliberately chose this value to make the problem of secondary maxima more pronounced in our example. If the likelihood transform approach can accelerate convergence in this kind of problem it will readily solve the same problems when the SNR is lower.

For this one-dimensional example, it is possible to calculate the likelihood on a grid of parameter values, using both the unsmoothed and smoothed waveform models. Figure 7 shows the computed likelihood for different smoothing widths. The desired smoothing properties are apparent: the global peak has broadened but remains close to the true value, and the number of secondary maxima has been reduced.

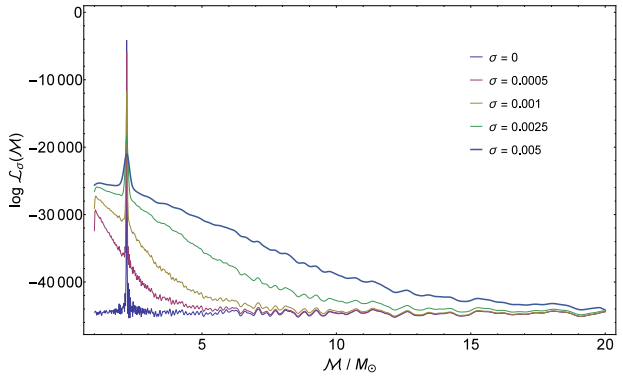

FIG. 7. The likelihood function computed using the smoothed waveform model for different values of the smoothing width. Thicker lines correspond to larger values of σ.

We now illustrate the more realistic situation of performing an MCMC search on the posterior distribution. We use a uniform prior on the chirp mass, and a Gaussian proposal distribution of width 10^{-3}. We choose 500 seeds $\mathcal{M}_{\rm seed}$ drawn from a uniform distribution across the allowed range of chirp masses, and start different MCMCs from each value: one using the unsmoothed surrogate waveform models; and four others using the smoothed model with different values of σ.

We run the chains for a small number of Metropolis-Hastings steps: 2500 for the unsmoothed chain and 1000 for the smoothed chains. We do not expect the chains to have converged on a stationary distribution by this point; it instead gives an indication of how quickly convergence may occur. The chain lengths are chosen such that the computational time is roughly equal for each type of chain.

The distributions of final chirp masses for both the unsmoothed and smoothed cases are shown in figure 8; the unsmoothed chains locate local maxima close to their seed value and so the final distribution is roughly uniform. On the other hand, the smoothed cases show that a significant number of chains have ended up near the global maximum at $\mathcal{M} = 2.2M_\odot$. The large peak visible at $\mathcal{M}_{\rm final} \approx 15M_\odot$ is a local maximum that has accrued many chains within the small number of steps that we have run; as the number of steps is increased we expect these chains to move towards the global peak, as illustrated by the chains in the vicinity of $5M_\odot$. Figure 9 shows the final chain values as a function of the seed chirp mass. Systems that start close to the global maximum locate it quickly using the smoothed waveform model.

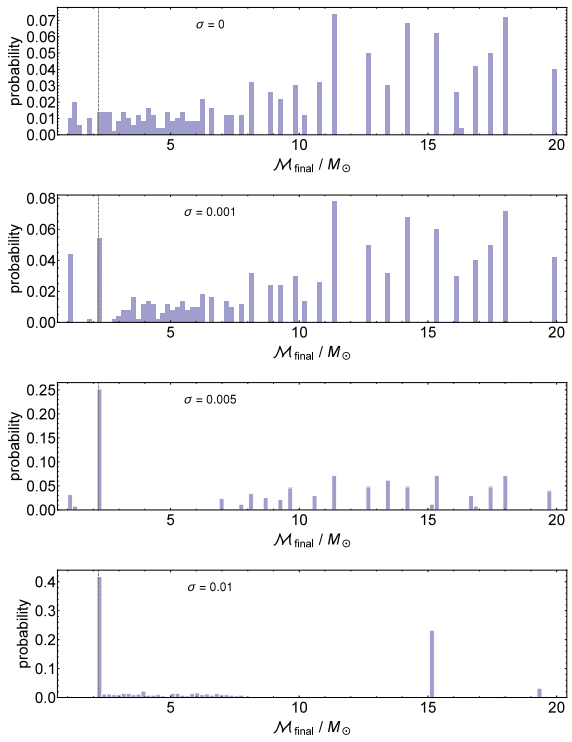

FIG. 8. Distribution of final chirp masses after a short MCMC, using likelihood functions with different amounts of smoothing. The narrowness of the peaks is indicative that the chains have located local maxima. The vertical grey dashed line is positioned at the true value $\mathcal{M}_* = 2.2M_\odot$.

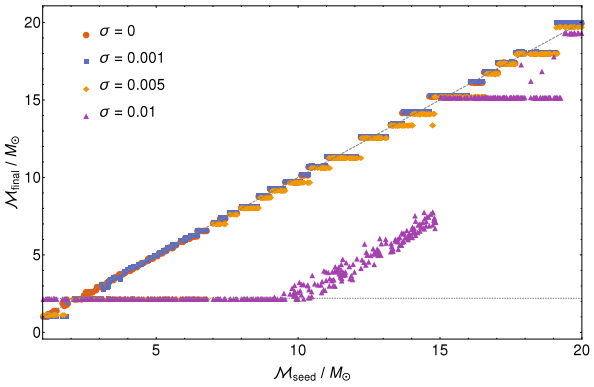

FIG. 9. The final chirp mass, as a function of the seed mass, for MCMCs using different smoothed likelihood functions. The dashed sloped line is $\mathcal{M}_{\rm final} = \mathcal{M}_{\rm seed}$, indicating chains that did not move far from their starting point. The dotted line $\mathcal{M}_{\rm final} = \mathcal{M}_*$ denotes the true value.

It is also possible to perform a comparison with simulated annealing and parallel tempering techniques. We run MCMCs for 2500 steps on the unsmoothed likelihood but at different temperatures, starting at the same chirp mass seed values as above. The final distributions of chirp masses are shown in figure 10. Changing the temperature of the chain does help to explore the parameter space, but is not as effective as likelihood smoothing at locating the largest peaks.

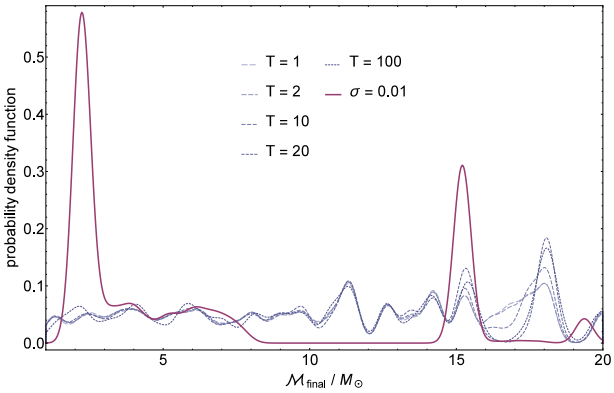

FIG. 10. Distribution of final chirp masses for chains at different temperatures (dashed lines), compared to the smoothed chain (solid line).

B. Higher order post-Newtonian chirp

We now consider our full PN waveform (up to 3.5PN in the phase and 2PN in the amplitude), with the two mass parameters $\mathcal{M}$ and ν allowed to vary. We look at chirp masses in the range $1 \leq \mathcal{M}/M_\odot \leq 20$ and mass ratios in the range[10] $0.1 \leq \nu \leq 0.25$.

1. *Idealised data*

For an initial study, we look at waveforms in the absence of noise. Our injected waveform has parameters $\mathcal{M}_* = 2.2M_\odot$, $\nu_* = 0.18$, $t_c = 0.1$s and $R = 1$Mpc, and is plotted in figure 11. The typical scales of variation about the true values obtained from the FIM are 2×10^{-3} for θ_1 (unit chirp mass) and 0.12 for θ_2 (unit symmetric mass ratio).

We construct a surrogate model for the two-dimensional PN waveform family using a training set of 190×190 waveforms, sampled at a frequency $f_{\rm sample} = 20\,$kHz. The chirp masses for the training set waveforms are the same as those in section IV A, while ν is selected

[10] If the component masses are $10M_\odot$ and $1.4M_\odot$, corresponding to a fiducial black hole-neutron star binary, the symmetric mass ratio is $\nu = 0.107725$.

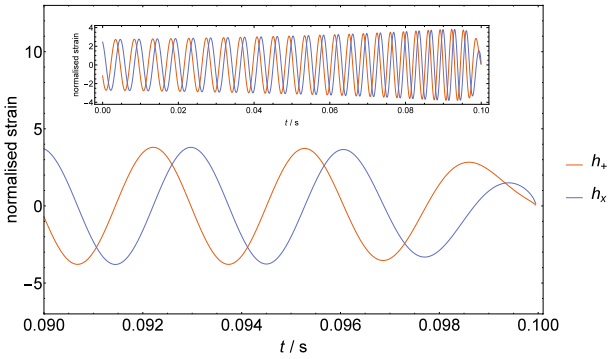

FIG. 11. PN gravitational waveform from a compact binary with chirp mass $\mathcal{M} = 2.2M_\odot$ and symmetric mass ratio $\nu = 0.18$ in the final 0.01s before coalescence. Inset is the same waveform for a period of 0.1s before coalescence. The amplitude is set by the requirement that $(h\,|\,h) = 1$.

uniformly from the allowed range for each value of $\mathcal{M}$. We target a representation error of 10^{-12} and the resulting RB contains 69 elements[11]. The error as a function of the size of the RB is plotted in figure 12. We confirm the faithfulness of the RB by computing the representation error for 2500 waveforms drawn from a uniform distribution on the parameter space; no error exceeded 10^{-12} and typical errors were $\mathcal{O}(10^{-14})$.

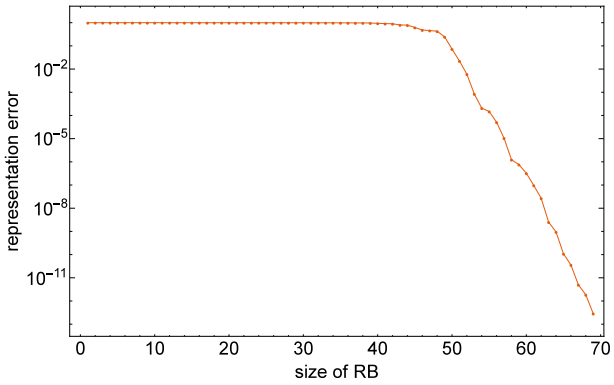

FIG. 12. The representation error of the RB for the PN waveform family as a function of the number of elements in the basis.

Following the surrogate model procedure, we use the training set waveforms identified in generating the RB to perform a fit to the waveform phase at each empirical node. In this work, we are not interested in the efficacy of surrogate models, but instead on their application to likelihood transform methods. To obtain an accurate surrogate model, we therefore use a fitting function based on the exact binary phase (A2). As a result, the fits are exact to within numerical precision.

[11] This number should not be directly compared to the 133 RB elements in section IV A since the waveform models are different.

The amplitude at each empirical node is now some complicated function of $\mathcal{M}$ and ν. To approximate this at arbitrary values, we perform plane interpolation between the greedy data points, which form an unstructured grid on parameter space. This seems somewhat less than ideal, however, the resulting surrogate model is sufficiently good due to our accurate phase model as well as the fact that the greedy points chosen by the RB algorithm are in important regions of parameter space.

We follow the approach of section III C and define a smoothed waveform model according to (44). We choose smoothing widths such that $\sigma_2 = 57\sigma_1$, consistent with the FIM estimates. The resulting smoothed waveforms for our injection signal for different values of $\sigma \equiv \sigma_1$ are shown in figure 13.

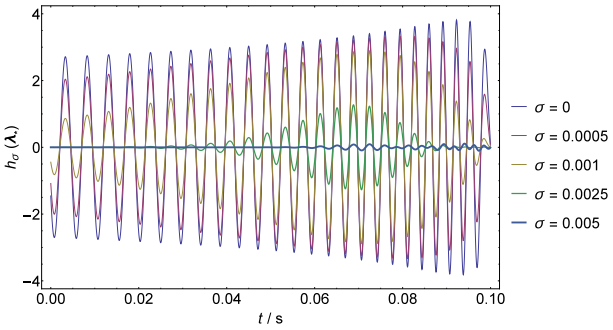

FIG. 13. The smoothed PN waveform, computed for our injection parameters with a chirp mass $\mathcal{M} = 2.2M_\odot$ and symmetric mass ratio $\nu = 0.18$. Thicker lines correspond to larger values of the smoothing width σ. The $\sigma = 0$ waveform is identical to that in figure 11.

The value of $\rho_\sigma(\boldsymbol{\lambda}_*; \boldsymbol{\lambda}_*)$ gives an indication of typical scales on which σ is likely to have the desired effect. Figure 14 shows the smoothed SNR as a function of σ; widths around 10^{-3} give a sufficient level of smoothing, consistent with the FIM. The strange behaviour at larger values of σ is due to the peak of the SNR distribution moving away from the true value $\boldsymbol{\lambda}_*$.

2. *Noisy data*

As with the 0PN waveform, we now consider the plus polarisation of our injection signal in the presence of white noise. With just two parameters, it is still feasible to map out the likelihood surface for different amounts of smoothing, shown in figure 15 along with an indication of the size of the smoothing region. As the value of σ is increased, we see a decrease in the amount of structure on the likelihood surface.

V. CONCLUSIONS

We have devised a practical scheme for using gravitational wave surrogate models to perform likelihood trans-

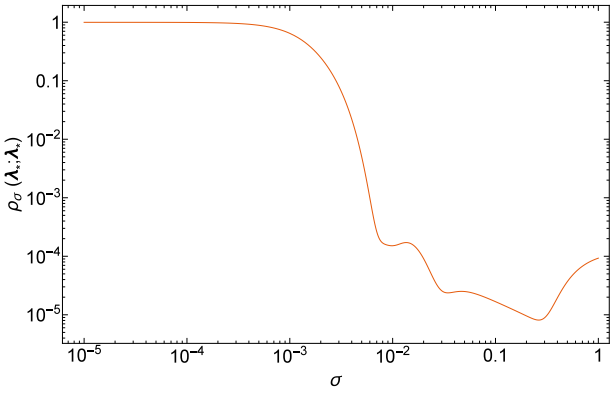

FIG. 14. The smoothed SNR for the PN waveform family, computed for our injection parameters with a chirp mass $\mathcal{M}_* = 2.2M_\odot$ and symmetric mass ratio $\nu_* = 0.18$, as a function of smoothing width.

form techniques, and have demonstrated that this can be used to accelerate convergence of MCMC methods. The advantage of smoothing the likelihood surface, rather than simply rescaling it (as is done in simulated annealing), is that the number of secondary maxima is reduced. The convolution required to perform this smoothing is expensive to evaluate numerically as it would require the generation of many waveform models. We make use of the waveform interpolation involved in constructing surrogate models to perform the convolution analytically.

Chirp waveforms have been considered as toy examples to demonstrate the methodology. We have calculated the overlap between, and corresponding likelihood of, model templates and injected data to illustrate the smoothing properties of the technique. We also considered an ensemble of short MCMCs in the presence of white noise, using unsmooothed and smoothed waveform models: the smoothing process accelerates the convergence of the algorithm without increasing computational time.

In practice, this approach could be implemented in a similar way to simulated annealing: starting the MCMC with a large smoothing width and gradually reducing this to zero according to some predetermined schedule. Alternatively, many different chains could be run with different smoothing widths, along with interchain communication, analagous to parallel tempering. A comparison of these different methods, using higher-dimensional waveform models, will be investigated in future work.

ACKNOWLEDGMENTS

RHC is supported by STFC. JG is supported by the Royal Society.

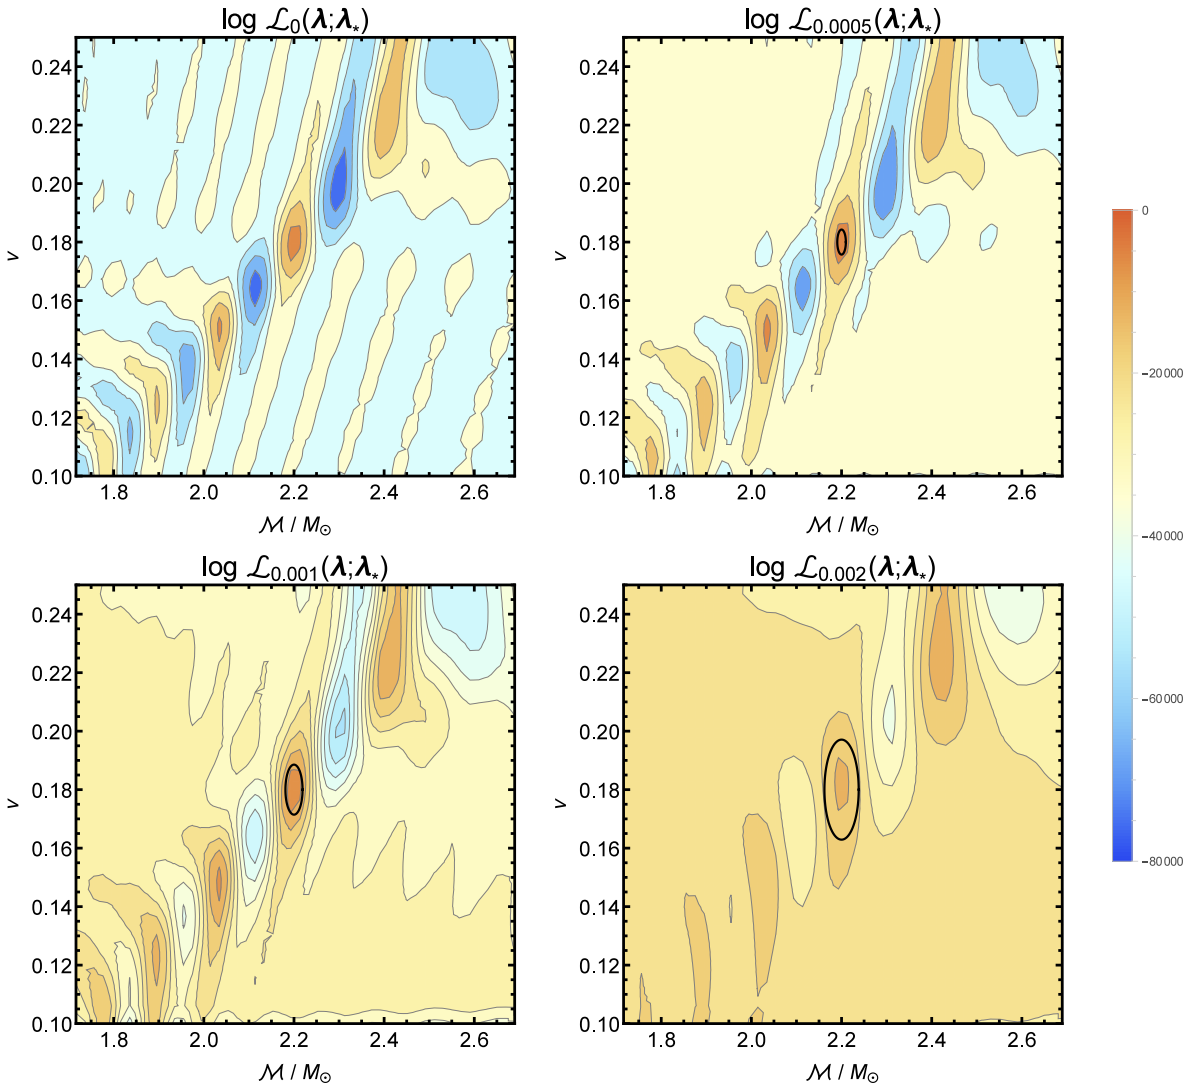

FIG. 15. The smoothed log-likelihoods for our PN waveform family using a selection of smoothing widths. The black ellipses are centred on the injection values and are sized according to the smoothing width.

[1] B. P. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. S. Amin, S. B. Anderson, W. G. Anderson, M. A. Arain, et al., Rep. Prog. Phys. **72**, 076901 (2009).

[2] T. Accadia, F. Acernese, M. Alshourbagy, P. Amico, F. Antonucci, S. Aoudia, N. Arnaud, C. Arnault, K. G. Arun, P. Astone, et al., J. Inst. **7**, P03012 (2012).

[3] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, P. Ajith, B. Allen, et al., Class. Quantum Grav. **27**, 173001 (2010).

[4] A. Sesana, MNRAS **433**, L1 (2013).

[5] P. Amaro-Seoane, S. Aoudia, S. Babak, P. Binétruy, E. Berti, A. Bohé, C. Caprini, M. Colpi, N. J. Cornish, K. Danzmann, et al., Class. Quantum Grav. **29**, 124016 (2012).

[6] C. J. Moore, R. H. Cole, and C. P. L. Berry, arXiv:1408.0740 (2014).

[7] G. M. Harry and the LIGO Scientific Collaboration, Class. Quantum Grav. **27**, 084006 (2010).

[8] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Science **220**, 671 (1983).

[9] Y. Wang, arXiv:1402.6211 (2014).

[10] P. Canizares, S. E. Field, J. R. Gair, and M. Tiglio, Phys. Rev. D **87**, 124005 (2013).

[11] M. Pürrer, Class. Quantum Grav. **31**, 195010 (2014).

[12] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, Phys. Rev. X **4**, 031006 (2014).

[13] L. Barack and C. Cutler, Phys. Rev. D **69**, 082005 (2004).

[14] S. Babak, H. Fang, J. R. Gair, K. Glampedakis, and S. A. Hughes, Phys. Rev. D **75**, 024005 (2007).

[15] L. Blanchet, Living Reviews in Relativity **17** (2014).

[16] L. Blanchet and G. Schafer, Class. Quantum Grav. **10**,

2699 (1993).
[17] L. Blanchet, G. Faye, B. R. Iyer, and S. Sinha, Class. Quantum Grav. **25**, 165003 (2008).
[18] L. Blanchet, B. R. Iyer, C. M. Will, and A. G. Wiseman, Class. Quantum Grav. **13**, 575 (1996).
[19] B. P. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. S. Amin, S. B. Anderson, W. G. Anderson, M. A. Arain, et al., Phys. Rev. D **79**, 122001 (2009).
[20] L. Blanchet, G. Faye, B. R. Iyer, and B. Joguet, Phys. Rev. D **65**, 061501 (2002).
[21] L. Blanchet, T. Damour, G. Esposito-Farèse, and B. R. Iyer, Phys. Rev. Lett. **93**, 091101 (2004).

Appendix A: Binary phase to 3.5PN order

To complete the description of the GWs emitted from a circular black hole binary used in section IV, we need an expression for the phase and frequency as a function of time. Using the auxilliary time variable

$$\tau(t) \equiv \frac{\nu c^3}{5GM}\left(t_c - t\right), \tag{A1}$$

where t_c is the coalescence time of the binary, the phase can be computed to 3.5PN order as [20, 21]

$$\begin{aligned}\Phi = -\frac{1}{\nu}\Bigg\{&\tau^{5/8} + \left(\frac{3715}{8064} + \frac{55}{96}\nu\right)\tau^{3/8} - \frac{3}{4}\pi\tau^{1/4} + \left(\frac{9275495}{14450688} + \frac{284875}{258048}\nu + \frac{1855}{2048}\nu^2\right)\tau^{1/8} \\ &+ \left(-\frac{38645}{172032} + \frac{65}{2048}\nu\right)\pi\ln\left(\frac{\tau}{\tau_0}\right) + \left(\frac{831032450749357}{57682522275840} - \frac{53}{40}\pi^2 - \frac{107}{56}\gamma_{\rm E} + \frac{107}{448}\ln\left(\frac{\tau}{256}\right)\right. \\ &\left. + \left[-\frac{126510089885}{4161798144} + \frac{2255}{2048}\pi^2\right]\nu + \frac{154565}{1835008}\nu^2 - \frac{1179625}{1769472}\nu^3\right)\tau^{-1/8} \\ &+ \left(\frac{188516689}{173408256} + \frac{488825}{516096}\nu - \frac{141769}{516096}\nu^2\right)\pi\tau^{-1/4}\Bigg\}.\end{aligned} \tag{A2}$$

The constant τ_0 sets the initial conditions when the binary passes some fiducial frequency; we choose $\tau_0 = \tau(0)$. $\gamma_{\rm E}$ is the Euler-Mascheroni constant.

The orbital frequency of the binary $\Omega = \dot{\Phi}$ can be computed by differentiating (A2) with respect to t.